

Making the Best of JPEG at Very High Compression Ratios: Rectangular Pixel Averaging for Mars Pathfinder

E. E. Majani, W. C. Dias
Jet Propulsion Laboratory
4800 Ok Grove Drive
Pasadena, CA 91109

November 29, 1994

1 Introduction

As more and more NASA missions are turning to image compression to maximize their data return at constrained bit rates, and very often adopting JPEG as the centerpiece of their image compression system, they are noticing one limitation of JPEG: its poor performance at very high compression ratios (typically 32 anti above).

While most mission scientists are interested in compression ratios from about 4-to-1 up to 1 G-to-1, there exist some applications in which very high compression ratios are desired, such as in the case of Mars Pathfinder, a mission to land a camera, rover and other instruments on the Red Planet in July 1997 [1, 2].

The most commonly studied use for compression algorithms on space missions is in handling science imaging data. There is much pressure to maximize the amount of information per bit sent from the spacecraft, from which data rates are limited. Let the problem is to preserve as much detail as possible about an imaging objective one cannot predict. The approach is always to maximize resolution and minimize information lost to compression.

However, there are engineering uses for image data. On Pathfinder these include assessment of lander condition and deployed airbags, and rover navigation. These problems are different from science imaging. One still wishes to maximize the information per bit, of course. However, unlike science scenes, images for engineering at full resolution may contain unnecessary information, and this fact can be used to advantage. One knows in advance the features one wishes to see to conduct the assessments, and their size and position (at least approximately).

If JPEG is used as the image compression algorithm for the camera, then the compression ratios attainable while preserving the information of interest (certain objects of certain sizes) are much lower than necessary, as we will find later. This is due to the fact that at high compression ratios, JPEG produces unacceptable artifacts, due to the limitation of the size of the Discrete Cosine Transform (DCT) to 8, for which no clever quantization or entropy coding can compensate.

One important observation is that for images, at high compression ratios, most of the high-frequency transform coefficients are quantized to zero, challenging the view that they should be computed at all, and suggesting that they simply should be dropped or removed.

Low-pass filtering of the full resolution image followed by subsampling is suggested as a way to accomplish just that. The scheme we propose consists in 4 sequential operations, the first two taking place in the camera, and the last two taking place on the ground:

1. low-pass filtering and downsampling (DF is the downsampling factor), horizontally and vertically;
2. JPEG compression in the compression ratio range at which JPEG is the most performant;
3. JPEG decompression;
4. expansion to the image's original size: upsampling by the same factors as in step 1, followed by interpolation.

The justification of step 1 is due to the fact that it is possible in principle to discard part of the information present in the full resolution image, as long as, for example, at least 2.5 pixels cover each feature needing to be seen.

There are two important refinements to the downsampling strategy for operations images, which are relevant to Pathfinder and potentially to any other mission with a lander camera. These are: (1) selecting the downsampling factor (DF) as a function of the known distance to the objective and its size, and (2) the use of different downsampling factors horizontally and vertically.

The first refinement is to maximize the DF based on the known distance to the objective of known size. The closer a feature, the greater the DF. Second, for maximum utility, downsampling must vary horizontally and vertically. Since the camera's mast is stationary on the lander, the ratio between resolution in elevation and azimuth depends on distance to a viewing Objective on the lander or on the terrain. Resolution in azimuth degrades more gradually than resolution in elevation with distance from the camera along the planetary surface. Therefore if one wishes to be able to detect only objects above a certain size in any dimension, such as rover obstacles above some threshold size, DFs are best computed as separate factors in elevation and azimuth as a function of distance. Often larger DFs are possible in azimuth.

One important consequence of downsampling is a savings in compression time, since the size of the original image is reduced. A second one is ease of command development. Assuming for simplicity that imaging objectives are on a plane perpendicular to the camera mast, the appropriate DFs are easily computed as a function of the size of the objective and distance to it. Further compression at "reasonable" compression ratios with equal treatment of the horizontal and vertical directions is then possible, with the assurance that irrelevant scene details will not be coded, and that coding artifacts will be acceptable.

In this paper, we propose to study efficient and appropriate schemes for steps 1 and 4, that are rarely studied in concert, assuming steps 2 and 3 are ignored. We derive optimal interpolation filters assuming unweighted pixel averaging is performed in step 1, and compare their performance to the widely used bilinear interpolation filters. Then we study the performance of the overall system (steps 1- 4), and show that bilinear interpolation filters are very close to optimal for the compression ratio range of interest, in the Mean Square Error (MSE) sense. We show that large compression ratios are achievable while at the same time preserving the relevant information for engineering uses.

2 Optimal Low-Pass and Interpolation Filters

One common approach in finding the best low-pass and interpolation filters to be used in steps 1 and 4, consists in the minimization of the MSE between the original and the reconstructed images. This is equivalent to the maximization of the Peak Signal-to-Noise Ratio (PSNR), defined as

$$PSNR = 20 \log_{10} \frac{255}{RMSE}$$

where RMSE is the square root of the MSE.

In this paper, we limit ourselves to the simplest low-pass filtering scheme known as unweighted averaging. The problem now consists in finding the optimal interpolation filters, given that the low-pass filtering is simple unweighted averaging. Since the answer may vary depending on whether steps 2 and 3 take place (or JPEG is used at very low compression ratios), we treat those two cases separately.

2.1 Without JPEG compression (steps 1 and 4)

For a given downsampling factor DF , the MSE minimization can easily be formulated in multirate filterbank terminology as in [3]. While in [3], the MSE is minimized for an impulse signal, experiments with images (as we shall see) show a good correlation in MSE performance when one minimizes the total reconstruction MSE of a step signal. In the case where $DF=2$, if $H_0(z)$ and $G_0(z)$ are the low-pass and interpolation filters, then the MSE to be minimized is the sum of the distortion error

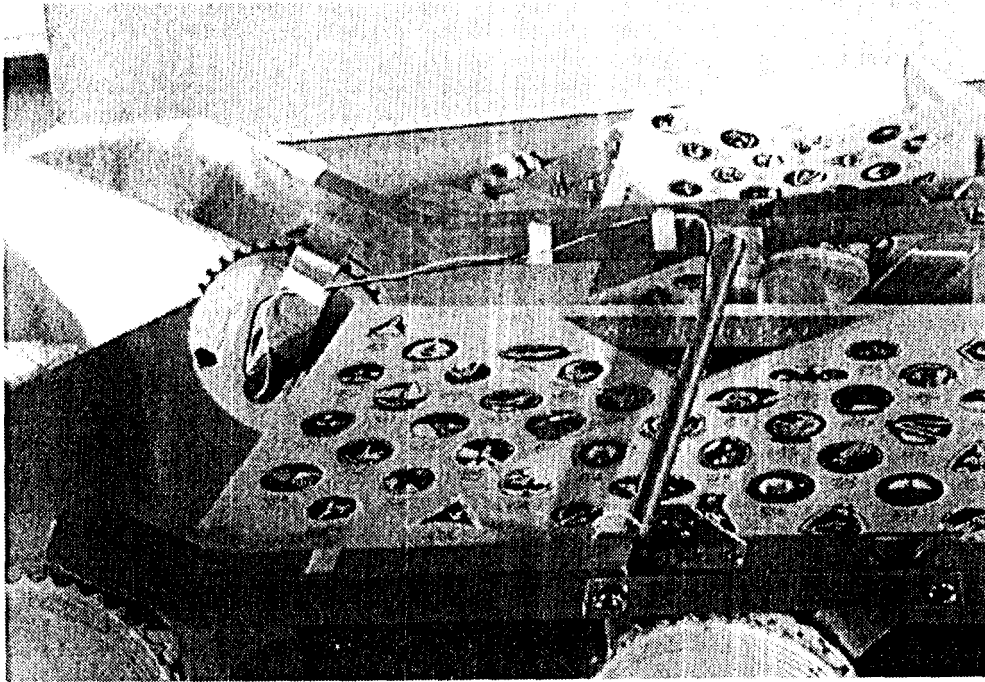
$$DE = \|(H_0(z)G_0(z) - 2)X(z)\|^2$$

and the aliasing error

$$AE = \|H_0(-z)G_0(z)X(-z)\|^2$$

introduced by the sampling process, where $X(z)$ is the Z-transform of a step signal.

Optimal interpolation filters G_0 can be found in this way for different filter lengths n , and they are given in Table 1, along with two PSNR values: PSNR1 corresponds to the PSNR obtained with a step signal, while PSNR2 corresponds to the PSNR obtained when applying steps 1 and 4 in the horizontal direction to a real image (Figure 1). Note that little PSNR improvement results from examining filters longer than $n=6$. Table 1 also contains the PSNR performance of the bilinear interpolation filter, the coefficients of which are $[1, 3, 3, 1]$ for $DF=2$. Its PSNR performance is about 1 dB below that of the best optimal filters. Note also the previously mentioned correlation between PSNR1 and PSNR2, i.e. interpolation filters optimized for a step signal perform well on images.



Original Image:480x336

Figure 1: Original Image

2.2 With JPEG compression (steps 1--4)

While the interpolation filter design technique can be predicted to yield good MSE performance when JPEG (steps 2 and 3) is used at low compression ratios, it is not clear that this good performance extends to the high compression ratio case. To answer that question, experimental rate-distortion curves derived from applying JPEG at various quality factors to the image in Figure 1 have been computed, first without any

n	G_0	PSNR1	PSNR2
2	[1, 1]	3.01	32.21
4	[1, 5, 5, 1]	3.80	33.18
6	[-1, 1, 7, 7, 1, -1]	4.47	34.15
8	[-5, -29, 35, 203, 203, 35, -29, -5]	4.49	34.21
∞		4.52	34.27
bilinear	[1, 3, 3, 1]	3.59	32.95

Table 1: Optimal Interpolation Filters G_0 of length n for $M = 2$ (bilinear interpolation is equivalent to the interpolation filter [1, 3, 3, 1]): PSNR1 is from the reconstruction error of a step signal of intensity 1, PSNR2 is from the reconstruction error of image 1 due to a horizontal downsampling factor of 2.

pixel averaging, then with pixel averaging in the horizontal direction only ($DF = 2$). The reference curve in Figure 2 corresponds to no pixel averaging. The “bilinear” curve corresponds to using pixel averaging and bilinear interpolation, while the “ $n=6$ ” curve corresponds to pixel averaging and interpolation with the optimal filter for $n = 6$. Note that while the optimal filter of length 6 noticeably outperforms the bilinear interpolator at the lower compression ratios as expected, the reverse is true for the compression ratio range of interest, i.e. the one corresponding to a PSNR improvement over no pixel averaging. As a consequence, for all values of the DF used in the following experiments, we will consider only bilinear interpolation filters, the coefficients of which are given in Table 2.

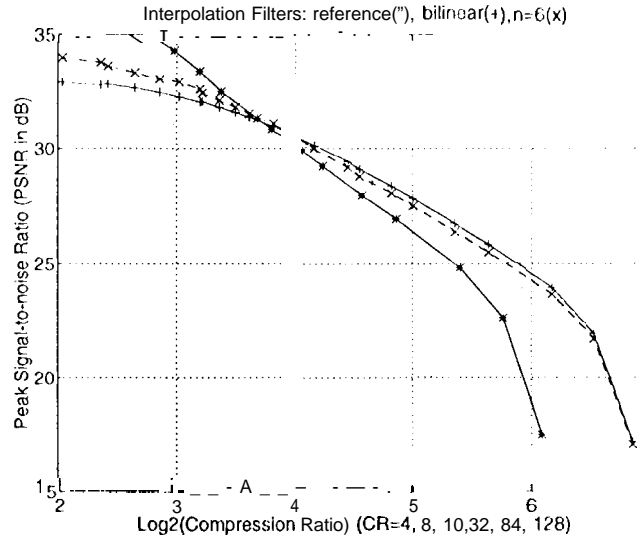


Figure 2: MSE (PSNR) Performance Comparison between Bilinear and Optimal Interpolation Filters combined with Lossy JPEG

DF	Filter Coefficients
2	[1, 3, 3, 1]
3	[1, 2, 3, 2, 1]
4	[1, 3, 5, 7, 7, 5, 3, 1]
6	[1, 13, 35, 7, 9, 11, 11, 7, 5, 3, 1]

Table 2: Bilinear Interpolation Filters for $DF = 2, 3, 4$ and 6

3 image description and Compression Assessment

The lander is in the shape of a tetrahedron. The initial landing on Mars is on 1 airbags. Airbag deployment problems or unexpected landing obstacle patterns could result in lander damage. After the landing, airbags are slowly deflated and retracted, an imprecise process which can also be affected by nearby rocks. The lander then unfolds the three motorized, hinged petals to coplanar. The battery holds only enough charge for a full day's normal operations, although this time could be stretched out to several days on an emergency basis. Therefore the conduct of the mission is dependent on the amount of solar power coming in through solar panels on the unfolded petals. This can be affected by twisting of the petals during rough landing, terrain slope, nearby obstacles or remaining folds in the airbag blocking the sun, or airbags or other obstacles slowing rover deployment. All the above conditions (except some kinds of lander damage) can be assessed with images compressed more than that acceptable for science.

The original image (Figure 1.) was acquired with a CCD camera with the optics chosen to closely mimic the real Pathfinder camera pixel size of 1 milliradian, and the real camera position about 60 cm high. The scene shown is of (1) a deployment test model of the rover in stowed position, (2) part of one of the lander's unfolded petals, (3) folded airbags visible around the left edge and tip of the petal, and (4) dark circular test patterns on the rover and near the petal tip, centaining fine detail unnecessary to the assessment task, about 2 cm in diameter.

Comparisons of decompressed images will now be made, without pixel averaging, and with three pixel averaging schemes. The notation $HmVn/p$ will refer to a horizontal DF of m , a vertical DF of n , and an overall compression ratio p . All H1V1 schemes therefore correspond to no pixel averaging.

3.1 Without Pixel Averaging (steps 2- -3)

In H1V1/17 (Figure 4a), we can see that the entire rover is adequately preserved for engineering assessment. The edge of the petal is clearly visible, and nearby airbag folds are seen to not shade the petals. The airbags near the tip of the petal have begun to be obscured, but even though they perhaps could not be clearly distinguished from terrain features, it is obvious they are small enough to not impede the solar input. Much detail unnecessary to lander assessment is preserved such as the light detail within the dark circular test patterns all the way out to the tip of the petal, and the bolts near the petal tip.

In H1V1/32 (Figure 4b), unnecessary high-frequency detail continues to be preserved. JPEG's blockiness has begun to obscure important features, though not fatally. One cannot determine whether the airbag near the tip blocks the sun's rays at some angles. Also partially obscured is whether the darker areas on the lander petals are shadows from something, reflections of the airbag, or perhaps even blown in surface material.

H1V1/62 (Figure 4c), is of such low quality that one would never plan the mission with the idea of relying on it. It is interesting to note, however, that unnecessary detail is still preserved, at least in a gross way, in the area of dark circular test patterns. While not a desirable image, a few features can be deduced or inferred; it is an image which could conceivably be used for something if it was all one had.

3.2 With Pixel Averaging (steps 1-4)

In Figure 3, we display the MSE performance of JPEG alone, as well as combined with pixel averaging on our sample image. Four pixel averaging schemes are considered: if m and n refer to the downsampling factors in the horizontal and vertical directions in the expression $HmVn$, then the schemes selected are no pixel averaging (H1V1), H2V1, H4V2, and H6V3. The axes of the graph are the logarithm base 2 of the compression ratio and the Peak Signal-to-Noise Ratio (PSNR) defined as

$$PSNR = 20 \lg 107 \left\{ \frac{1}{N} \sum_{i,j} \left(\frac{R_{ij}}{F_{ij}} \right)^2 \right\}$$

where R_{ij} denotes the square root of the mean square error between the original and decompressed images

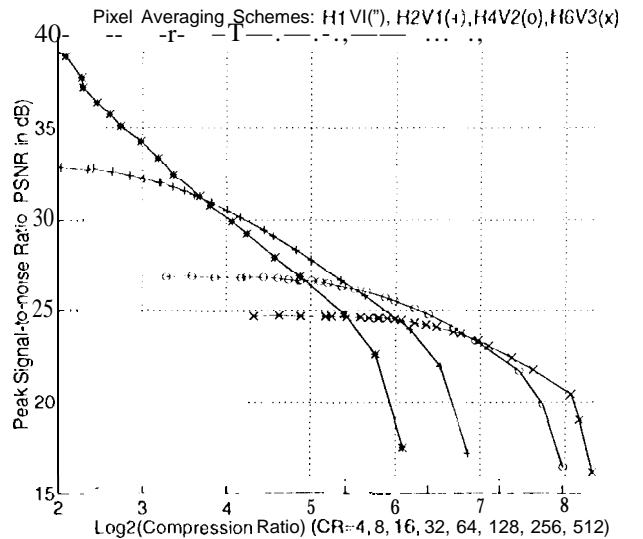


Figure 3: MSE (PSNR) Performance Comparison between Pixel Averaging Schemes combined with Public Domain JPEG: H_mV_n refers to downsampling factors of m horizontally and n vertically

$H_2V_1/17$ (Figure 5a) appears superior to $H_1V_1/17$. Even though it uses the same number of bits, due to its lesser blockiness, it is clearer that the dark areas in the airbags are shadows from airbag folds. It is also clearer that dark areas on the petal surface are reflections. High frequency detail within the dark circles has been sacrificed to achieve this, which is the correct priority.

$H_2V_1/32$ (Figure 5b) is superior to $H_1V_1/32$ for engineering purposes, and in fact appears to be as good as $H_1V_1/17$ for about half as many bits. Comments from $H_1V_1/17$ therefore apply. Additional unnecessary detail has been downsampled out of the dark circular areas and the unneeded bolts near the tip of the petal are washed out almost completely.

$H_4V_2/63$ (Figure 5c) is, while borderline as an engineering image at this distance from most of the objectives, clearly preferable to $H_1V_1/62$. With regard to the closest part of the image, the rover and all its parts are generally visible, however the lower left corner has become indistinguishable from the petal material. It is clear by the shadowing how the airbags on the left are folded, and that they do not shade the solar panels. Other airbag fold patterns might be harder to interpret for solar panel shading, and the image would be inadequate for assessing the ability of the rover to traverse through the folds. The most distant part of the airbag cannot be assessed. This level of compression would probably suffice for objectives closer to the camera than the inner edge of the rover, and might be adequate out to the middle field for assessment of shading from airbags.

$H_6V_3/126$, (Figure 6) while clearly below the quality one would aim for, is still usable for some purposes, and still much better than $H_1V_1/62$. The airbags on the left have blurred into the surrounding area outside the lander. Large rover features such as wheels have their basic shapes obscured. One still may be able to tell if the lander petals are twisted, and determine if some of the airbags shade solar panels. This level of compression could probably be used only for assessments very close to the camera, though even at this high compression ratio one can observe that JPEG is still preserving some of the unnecessary detail inside the 2cm dark circles. This suggests that adjustment of JPEG's quantization table to favor lower-frequency objectives could improve the overall usefulness of this and other images in the Lest group.

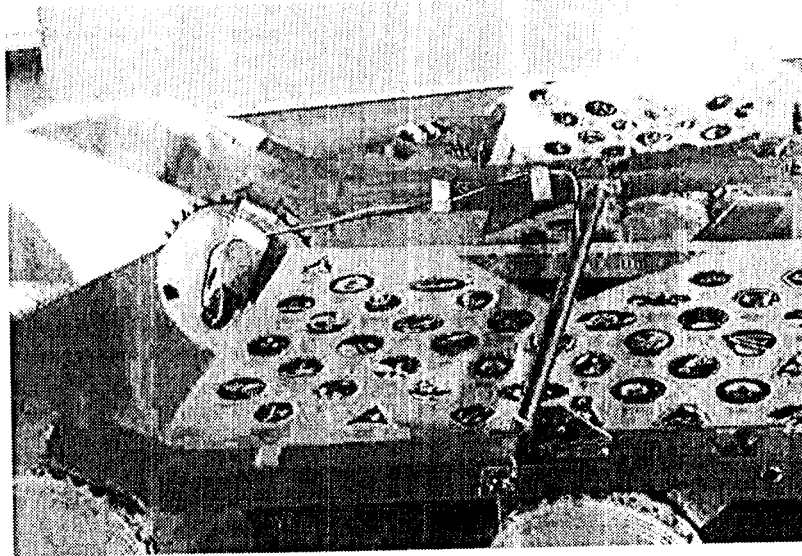
4 conclusion

Rectangular pixel averaging appears to be a useful form of image compression for engineering assessments on landed planetary missions. For Pathfinder this advantage is available at least up to about 8:1 and maybe 18:1 for some purposes. When combined with more intelligent image compression, such as JPEG, it provides the capability to provide adequate image quality for engineering assessment while greatly reducing the number of hits required.

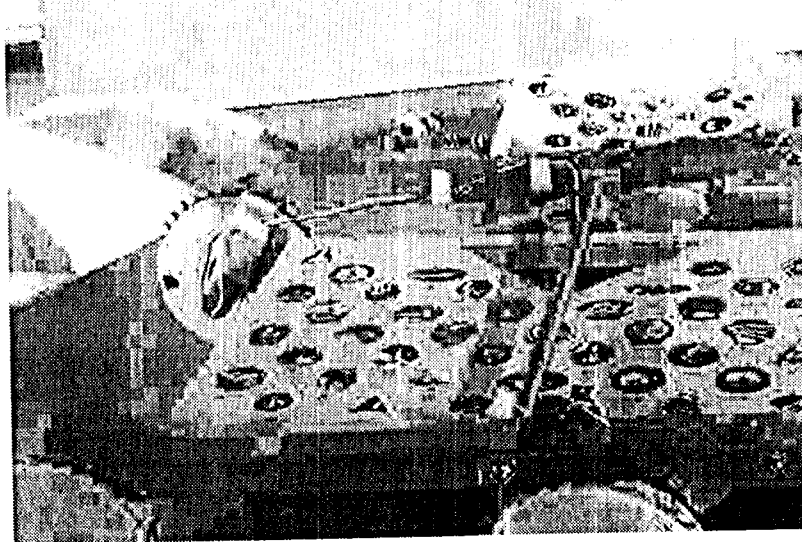
More generally, the use of pixel averaging as a pre-processing step to JPEG enhances the performance of JPEG at high compression ratios in the MSE sense, along with bilinear interpolation after JPEG decompression, as made clear from the experimental rate-distortion curves shown in this application.

References

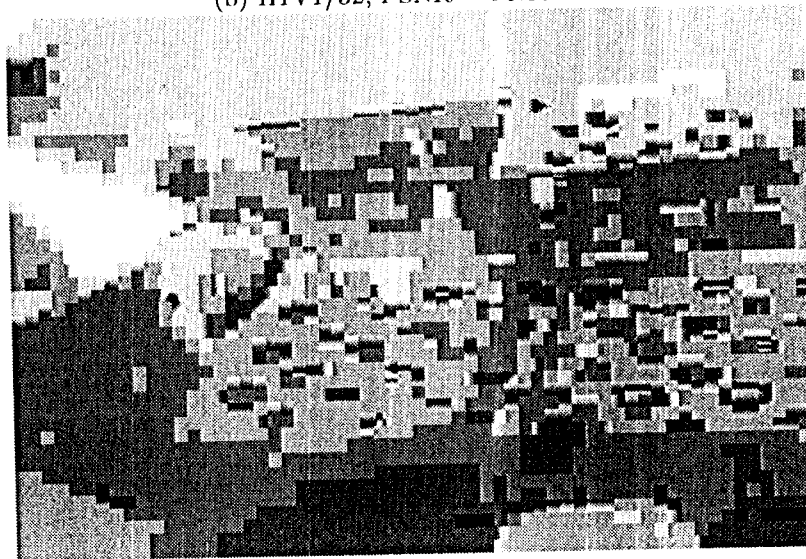
- [1] Anthony Spear, "Low Cost Approach to Mars Pathfinder and Small Landers", to appear in *Acts Astronautica*, Vol 35, Suppl., pp 345-354, 1995. Pergamon reprint 0094-5765(94)00200-2.
- [2] Anthony J. Spear, "Low Cost Approach to Mars Pathfinder", 45th Congress of the International Astronautical Federation, Oct 9-14, 1994, Jerusalem, Israel. 1 AF-94-Q.3.339.
- [3] M. Unser, M. Eden, "FIR approximations of inverse filters and perfect reconstruction filter banks" *Signal Processing*, vol. 36, pp. 163-174, 1994.



(a) H1V1/17, PSNR = 29.93 dB

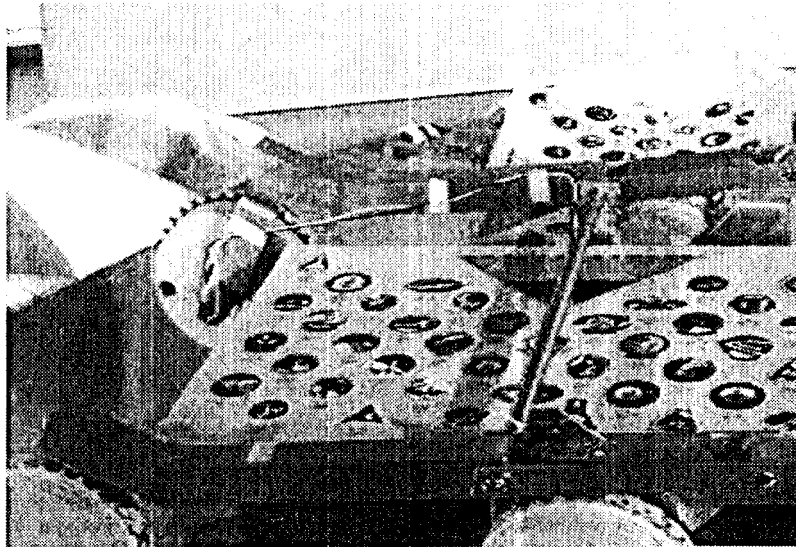


(b) H1V1/32, PSNR = 26.30 dB

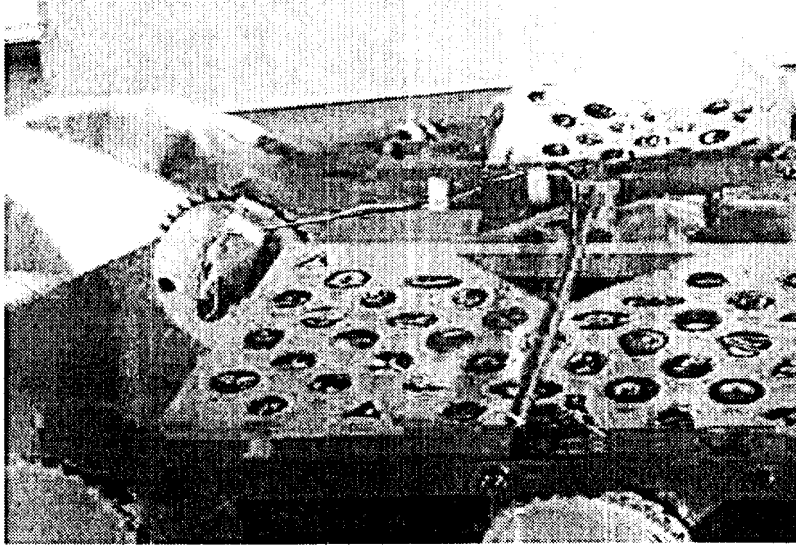


(c) H1V1/62, PSNR = 20.86 dB

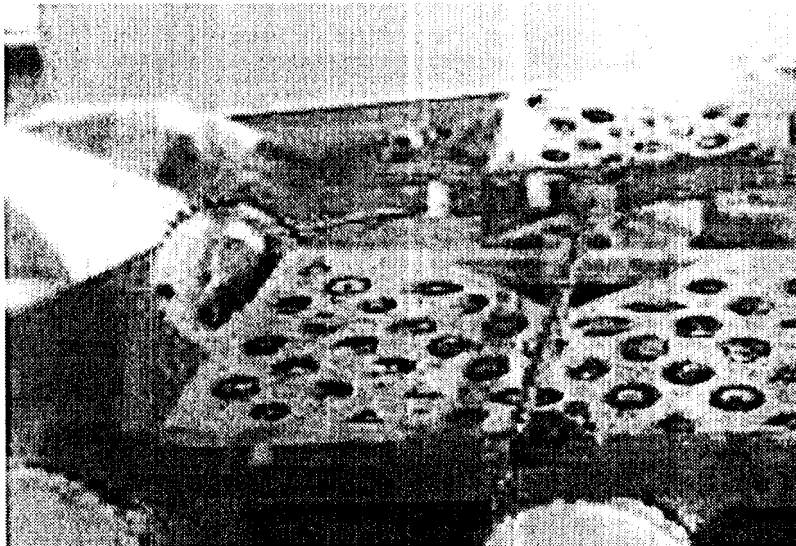
Figure 4: Decompressed Images



(a) H2V1/17, PSNR = 30.34 dB

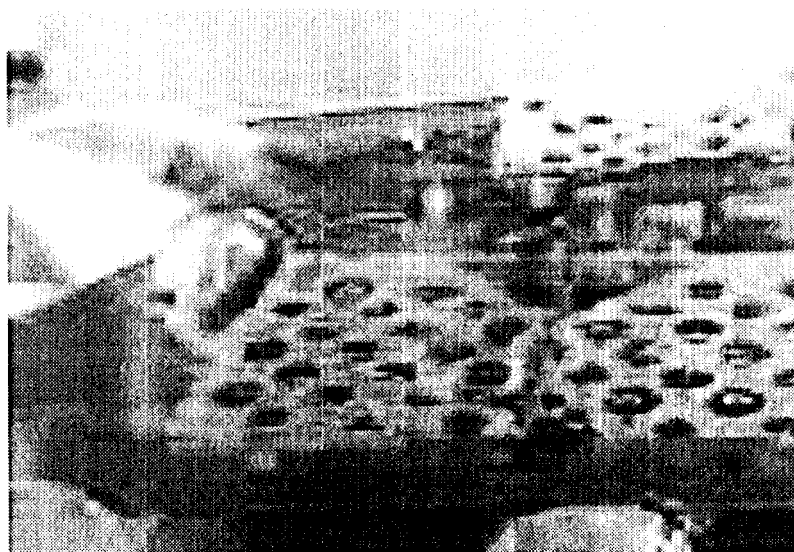


(b) H2V1/32, PSNR = 27.82 dB



(c) H4V2/63, PSNR = 25.56 dB

Figure 5: Decompressed images



H6V3/126, PSNR = 23.37 dB

Figure 6: Decompressed 1 images